

What is claimed is:

1 1. A method for use with a stateful packet processing device of a computer
2 network for mitigating effects of a network overload against said device, said method
3 operable to free memory used to store information about communications sessions
4 managed by said device, said method comprising the steps of:

5 classifying session cache entries made in memory into different cache classes,
6 according to one or more characteristics of those entries;

7 determining when said device is under network overload;

8 selecting session cache entries for deletion and deleting them thereby freeing
9 associated memory when said device is under network overload;

10 determining when sufficient memory has been freed, such that said cache entries
11 are no longer deleted.

1 2. The method of Claim 1, wherein said characteristics for said step of
2 classifying are selected from the group consisting of: whether the session is dropped by
3 the device, whether the session is audited by the device, IP protocol of the session, ICMP
4 type and code used in the session, TCP ports used in the session, UDP ports used in the
5 session, and whether the session is a half-open TCP session.

1 3. The method of Claim 1, wherein certain of said characteristics of the
2 session may be identified as "any", wherein any session matches a particular criterion for
3 classification.

1 4. The method of Claim 1, wherein predefined cache classes are selected
2 from the group consisting of:

3 dropped and unaudited sessions, dropped and audited sessions, ICMP sessions,
4 and half-open TCP sessions.

1 5. The method of claim 4, wherein the predefined cache classes are assigned
2 a priority for deletion.

1 6. The method of Claim 1, wherein the device is considered to be under
2 network overload when the amount of memory used for session cache entries exceeds a
3 configurable trigger threshold.

1 7. The method of Claim 6, wherein a sufficient amount of memory has been
2 freed when the amount of memory used for session cache entries falls below a
3 configurable floor threshold.

1 8. The method of Claim 4, wherein a memory usage threshold is
2 configurable for each predefined cache class.

1 9. The method of Claim 8, wherein said step of selecting and deleting
2 includes the steps of:

3 retrieving from a database the amount of memory used to store session cache
4 entries for each cache class ;

5 recognizing each cache class whose memory usage exceeds an associated memory
6 usage threshold;

7 ordering each cache class according to its deletion priority;

8 selecting for deletion according to said ordering step some fraction of entries of a
9 given cache class if said deletion brings said total cache memory usage below said floor,
10 wherein, otherwise, all entries of said given class are selected for deletion; and

11 continuing said step of selecting for deletion until it is determined that either
12 deleting all the entries selected for deletion would bring the total cache memory usage
13 below the floor threshold, or all entries in all defined cache classes have been selected for
14 deletion.

1 10. The method of Claim 9, wherein said step of ordering includes ordering
2 cache classes whose memory usage does not exceed said associated memory usage
3 threshold.

1 11. The method of Claim 9, wherein configuration data for the thresholds may
2 be supplied in a normalized fashion and be adaptively applied to the device, depending on
3 the amount of memory on the device.

1 12. An apparatus for use with a stateful packet processing device of a
2 computer network for mitigating effects of a network overload against said device, said
3 apparatus operable to free memory used to store information about communications
4 sessions managed by said device, said system comprising:

5 a classification component operable to determine, for each session cache entry, the
6 cache class to which that entry belongs according to one or more characteristics of the
7 entry;

8 a memory management database for tracking the amounts of memory used for
9 each category of entry, as well as for tracking the total amount of memory used for all
10 entries;

11 a pruning component that is used to select and delete entries; and

12 a processor for determining when said device is experiencing network overload
13 and selecting specific cache session entries for deletion until sufficient memory has been
14 freed.

1 13. The apparatus of Claim 12, wherein information kept in the memory
2 management database is updated each time a new cache entry is created or deleted by the
3 device.

1 14. The apparatus of Claim 12, wherein said characteristics for said step of
2 classifying are selected from the group consisting of: whether the session is dropped by
3 the device, whether the session is audited by the device, IP protocol of the session, ICMP
4 type and code used in the session, TCP ports used in the session, UDP ports used in the
5 session, and whether the session is a half-open TCP session.

1 15. The apparatus of Claim 14, wherein certain of said characteristics of the
2 session may be identified as "any", wherein any session matches a particular criterion for
3 classification.

1 16. The apparatus of Claim 12, wherein predefined cache classes are selected
2 from the group consisting of:

3 dropped and unaudited sessions, dropped and audited sessions, ICMP sessions,
4 and half-open TCP sessions.

1 17. The apparatus of claim 16, wherein the predefined cache classes are
2 assigned a priority for deletion.

1 18. The apparatus of Claim 16, wherein a memory usage threshold is
2 configurable for each predefined cache class.

1 19. The apparatus of Claim 12, wherein the pruning mechanism selects entries
2 for deletion by:

3 retrieving from a database the amount of memory used to store session cache
4 entries for each cache class ;

5 recognizing each cache class whose memory usage exceeds an associated memory
6 usage threshold;

7 ordering each cache class according to its deletion priority;

8 selecting for deletion according to said ordering step some fraction of entries of a
9 given cache class if said deletion brings said total cache memory usage below a floor
10 threshold, wherein, otherwise, all entries of said given class are selected for deletion; and

11 continuing said step of selecting for deletion until it is determined that either
12 deleting all the entries selected for deletion would bring the total cache memory usage
13 below the floor threshold, or all entries in all defined cache classes have been selected for
14 deletion.

21. The apparatus of Claim 19, wherein the pruning mechanism operates by making only one pass through a list of session cache entries in said device.

23. The system of Claim 12, wherein the memory usage statistics are collected using the Simple Network Management Protocol (SNMP).

24. The apparatus of Claim 12, wherein the pruning mechanism, when it has to delete some fraction of the entries in a given cache class, approximates the fraction b/t (where b is the total number of bytes of memory that must be freed and t is the total number of bytes of memory used to hold session cache entries for that cache class) with another fraction p/q , where $p \geq 1$ and q is likely to be small relative to the total number of cache entries in that class; and then frees p entries out of every q entries in that cache class on the list of session cache entries.

1 25. A cache management system used in connection with session-type packet
2 processing devices of a computer network, said system comprising:

3 a memory management database for storing communication traffic classification
4 and memory threshold values;

5 a memory monitor for tracking overall memory usage and determining when said
6 memory threshold values stored in said memory management database are reached;

7 a cache classifier used to determine a class into which a given session of
8 communications traffic falls; and

9 a pruner mechanism for selecting and pruning selected sessions of said packet
10 processing device in accordance with said communication traffic classification and
11 memory thresholds programmed in said memory management database when said
12 memory threshold value is reached.

1 26. The system of Claim 25 wherein said prune selector is operable to
2 selectively prune sessions of an ordered overlimit class if the memory used by said class
3 is greater than the difference between a global ceiling threshold and a global floor
4 threshold.

1 27. The system of Claim 26, wherein said prune selector is operable to prune
2 all sessions of said overlimit class if the memory used by said class is less than the
3 difference between said global ceiling threshold and said global floor threshold.

1 28. The system of Claim 27, wherein a next highest priority class is examined
2 to determine if memory used by said class is greater than a remaining difference between
3 said global ceiling threshold and said global floor threshold, said next highest priority
4 class being selectively pruned if said difference is greater than said remaining difference.

1 29. The system of Claim 28, wherein said prune selector is operable to prune
2 all sessions of said next highest priority class if the memory used by said class is less than
3 said remaining difference.

1 30. The system of Claim 25, wherein said devices are selected form the group
2 consisting of: network firewalls, routers, switches and hosts.